UNIVERSITY OF TEXAS ARLINGTON

Towards Predicting Efficient and Anonymous Tor Circuits

Armon Barton[1], Mohsen Imani[1], Jiang Ming[1], and Matthew Wright[2]

[1] University of Texas at Arlington, [2] Rochester Institute of Technology

Anonymity with Tor

- Millions of users
- 7,000 + Tor relays

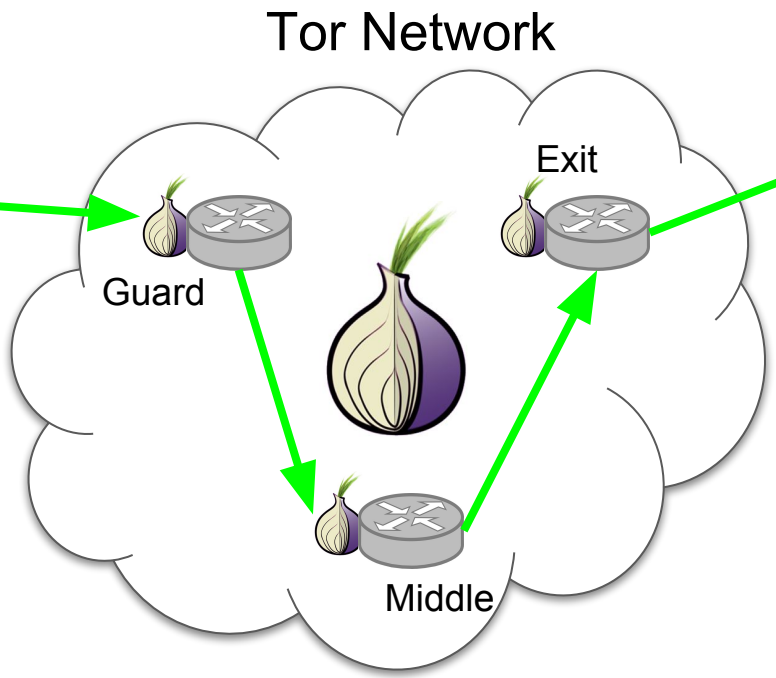

Anonymity with Tor

- Millions of users
- 7,000 + Tor relays

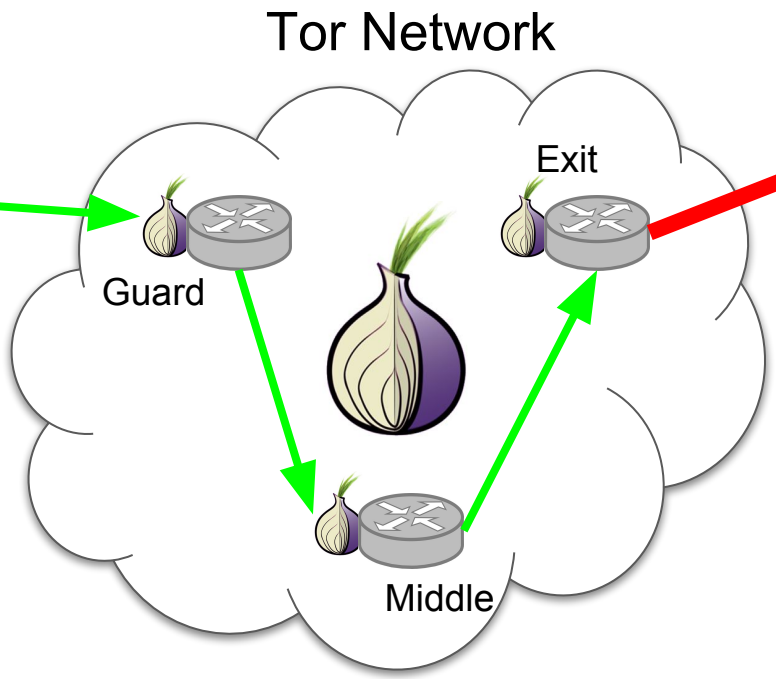

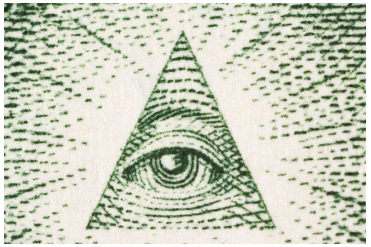

Anonymity with Tor

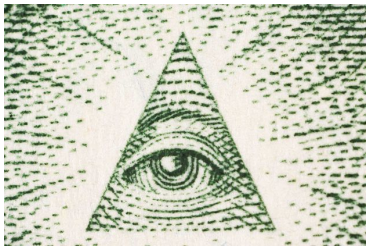

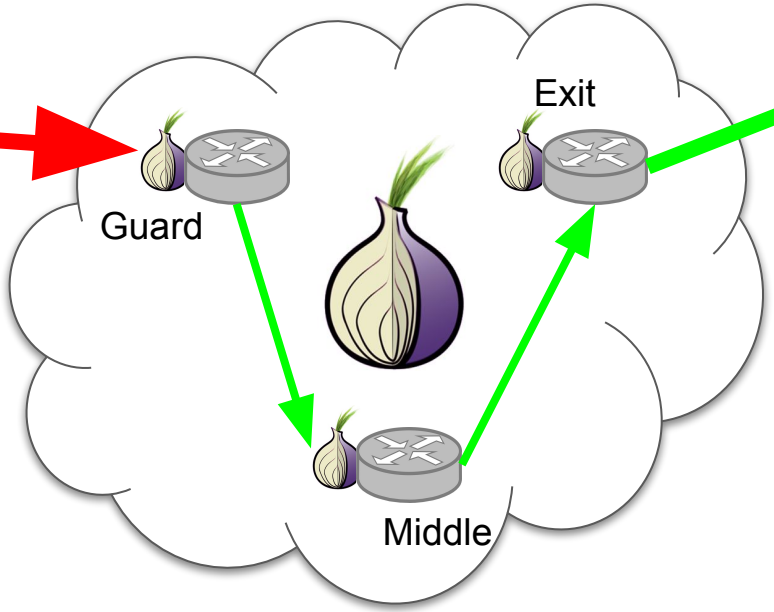

Latency in Tor

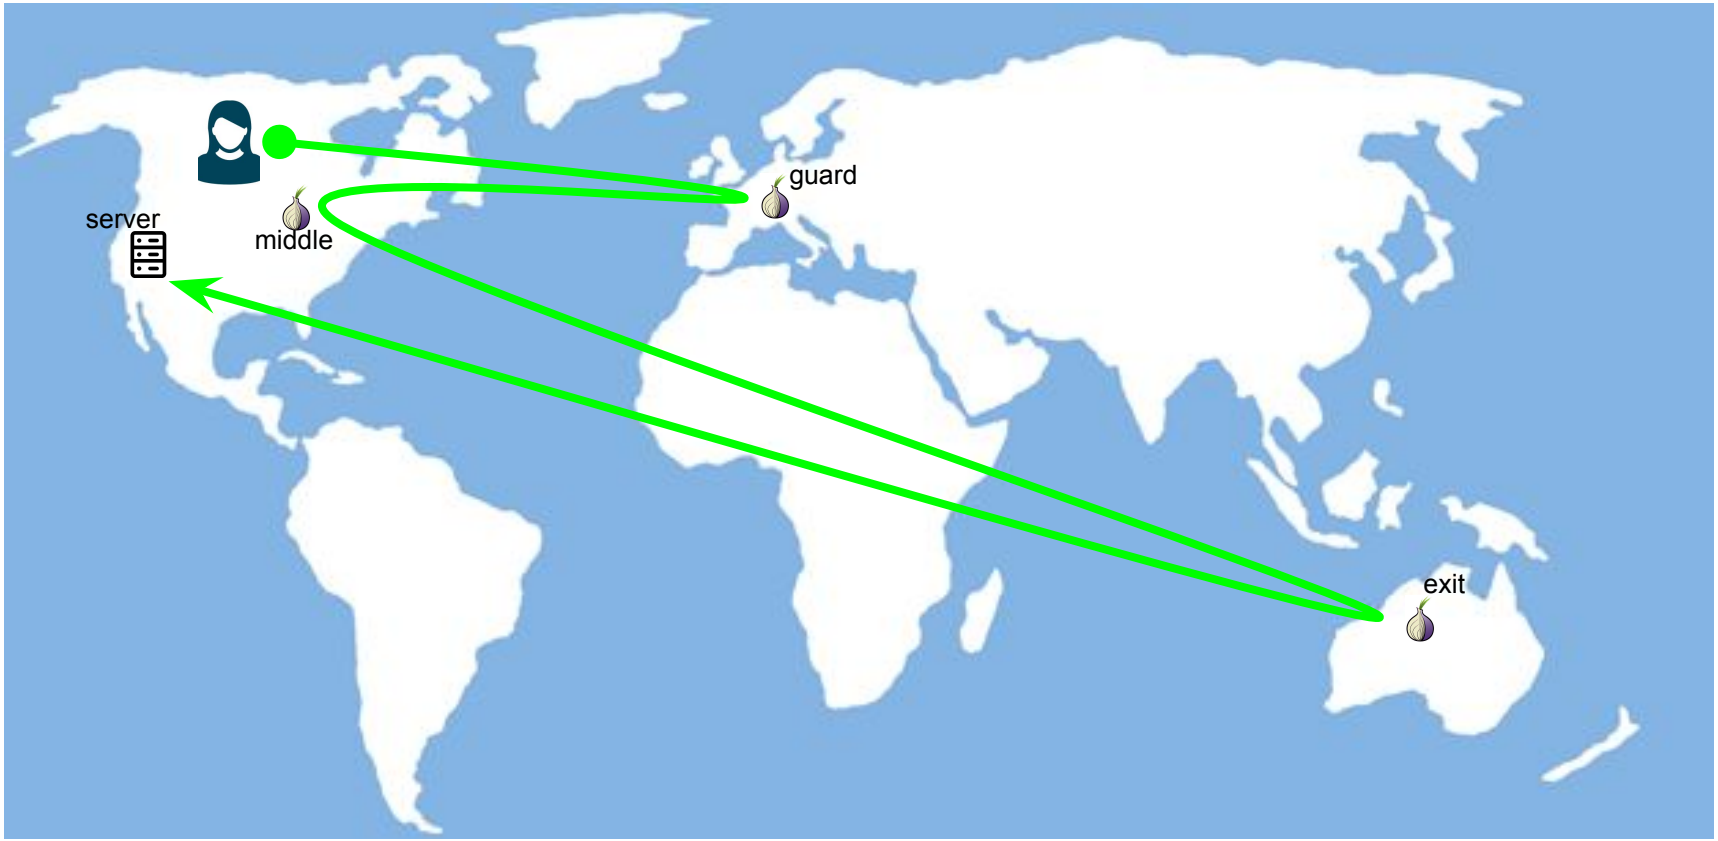

Tor Congestion

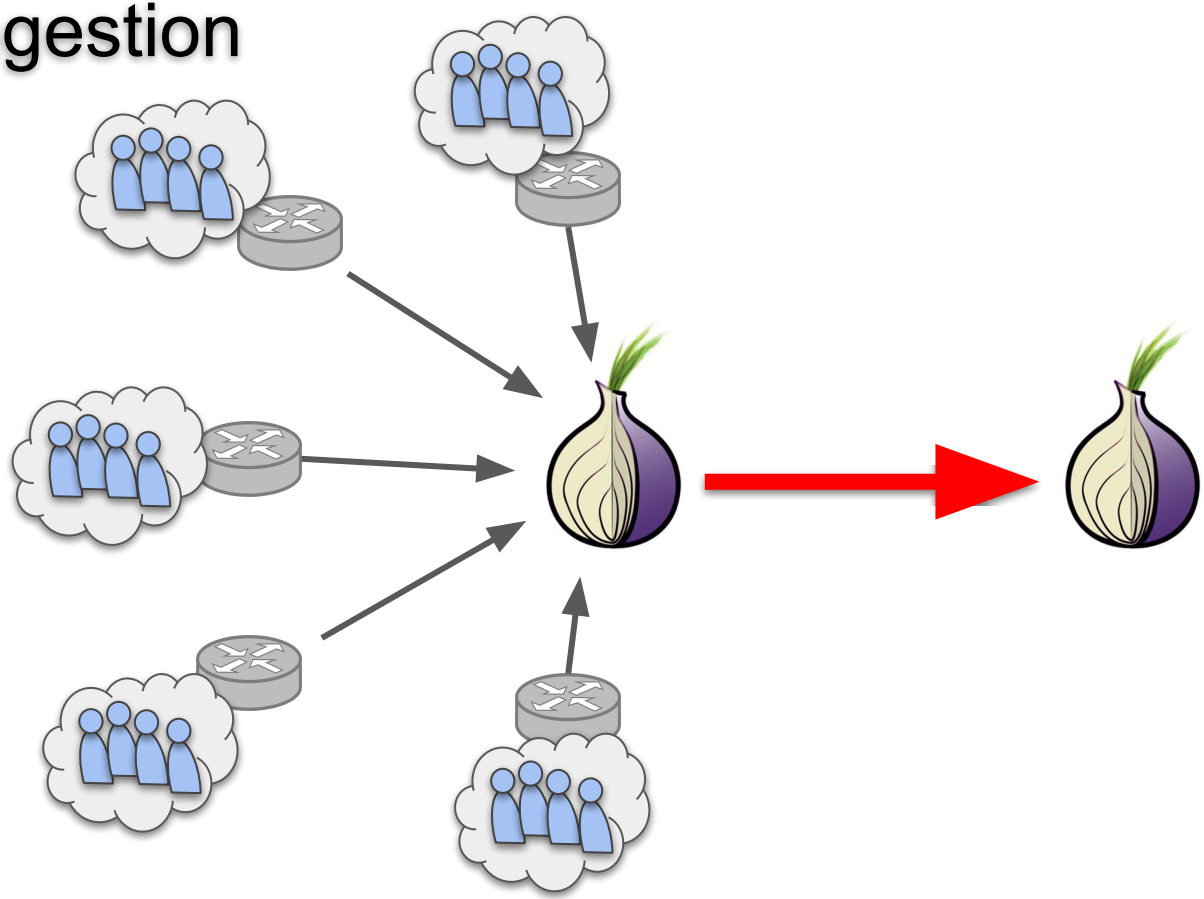

Path Selection Algorithms

- Bandwidth Weighted Selection
- Snader and Borisov Selection
- Congestion Aware Routing

Bandwidth Weighted Path Selection

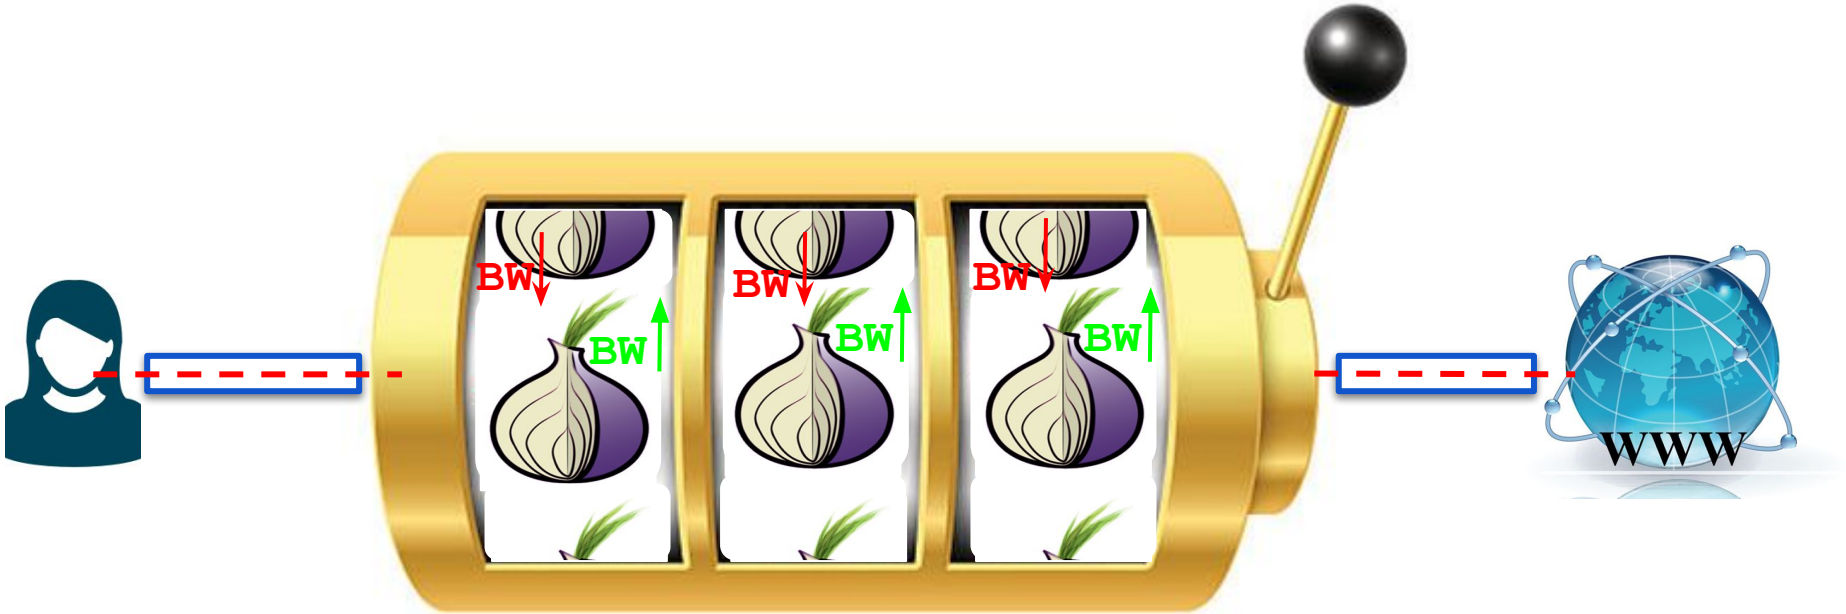

Tor Congestion

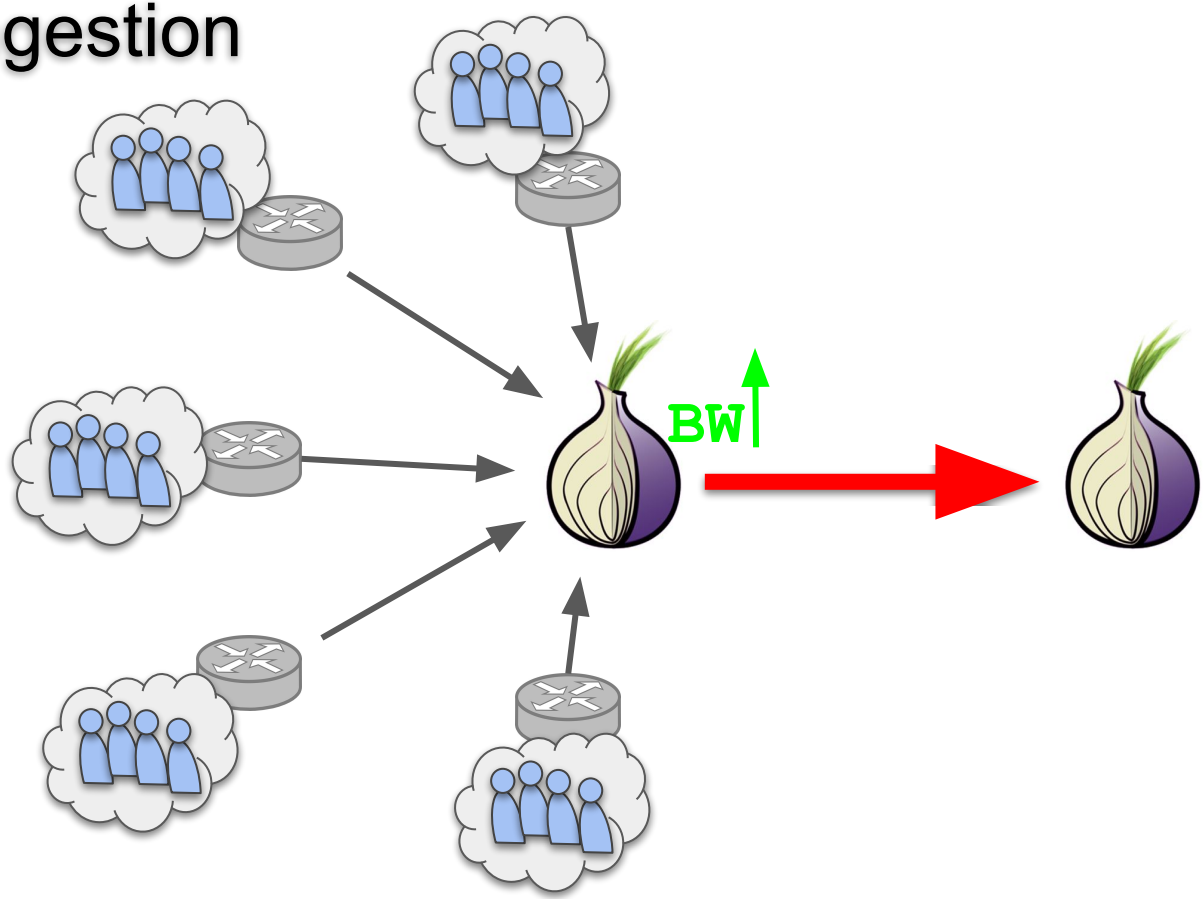

Snader and Borisov

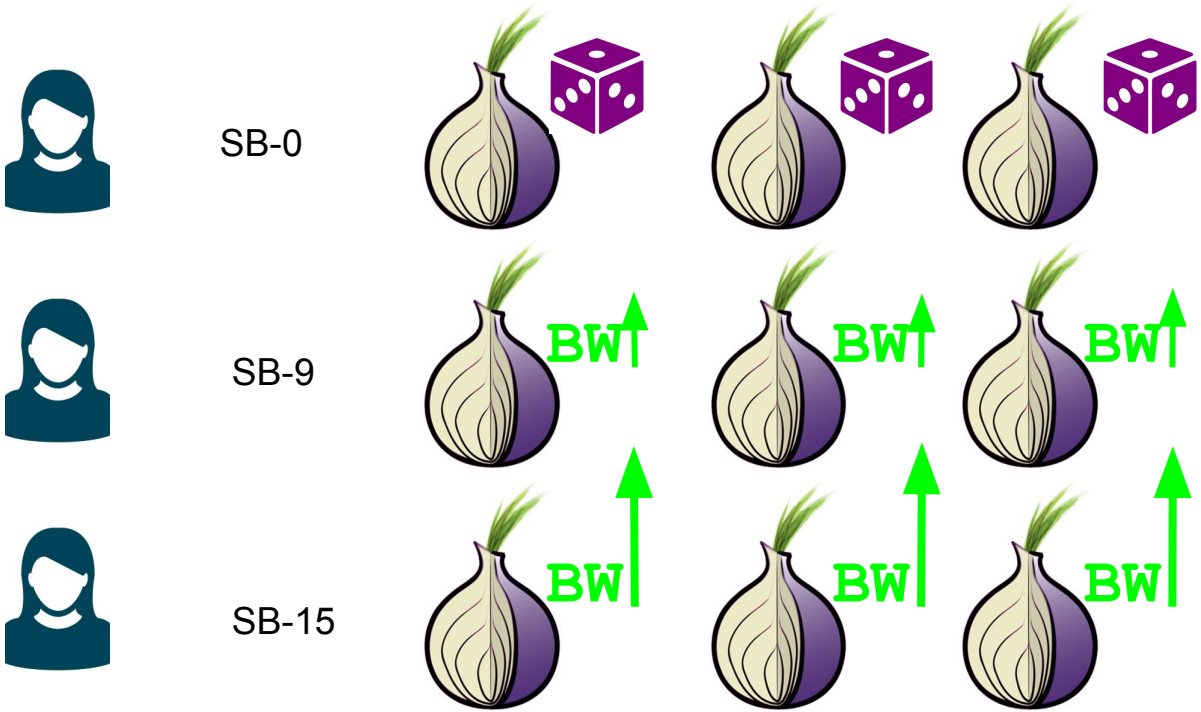

Tor Congestion

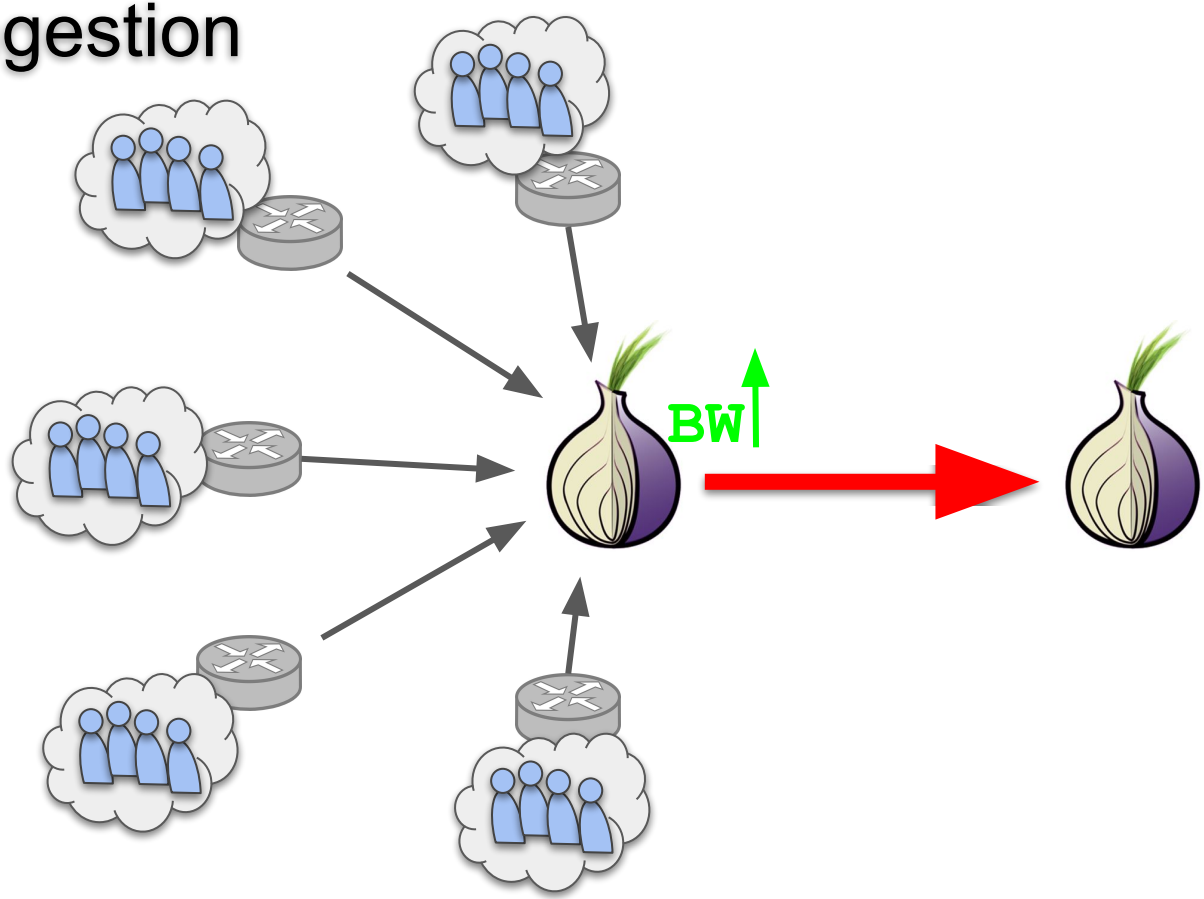

Snader and Borisov

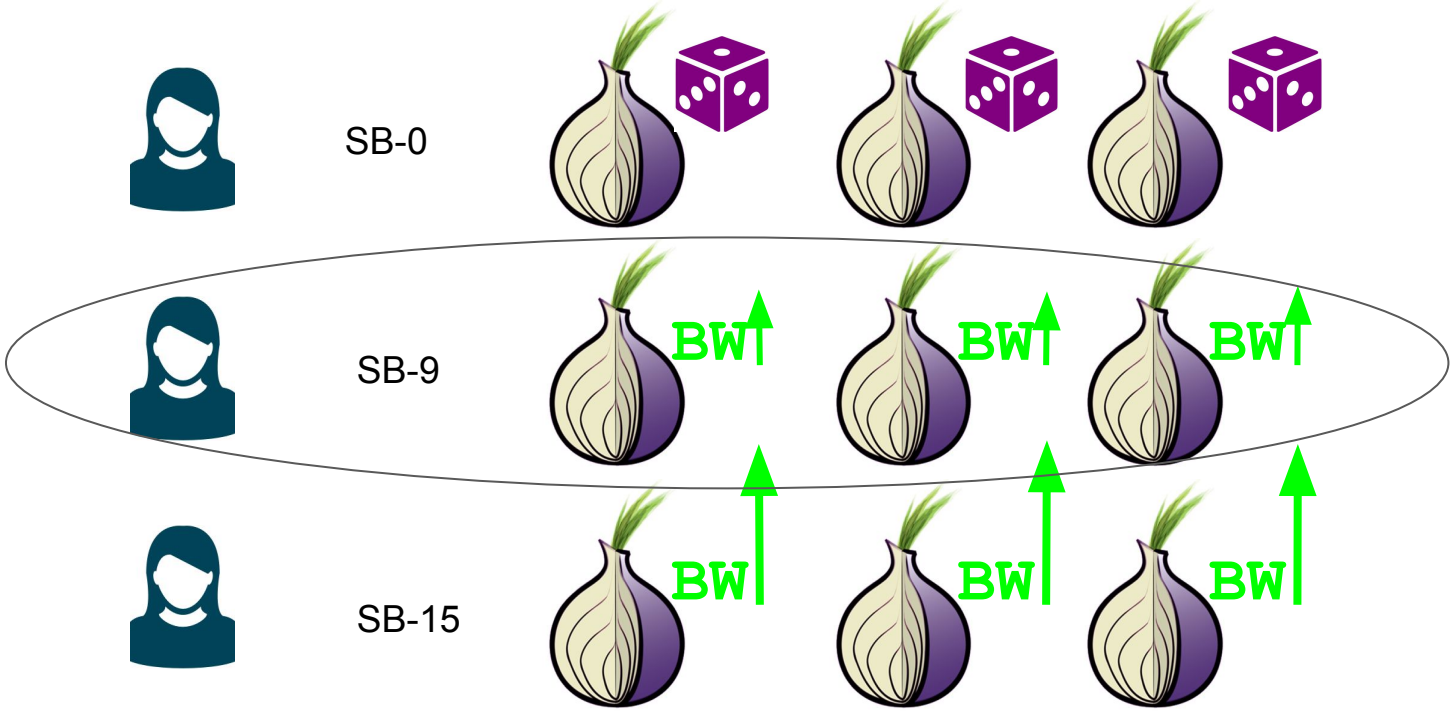

Congestion Aware Routing

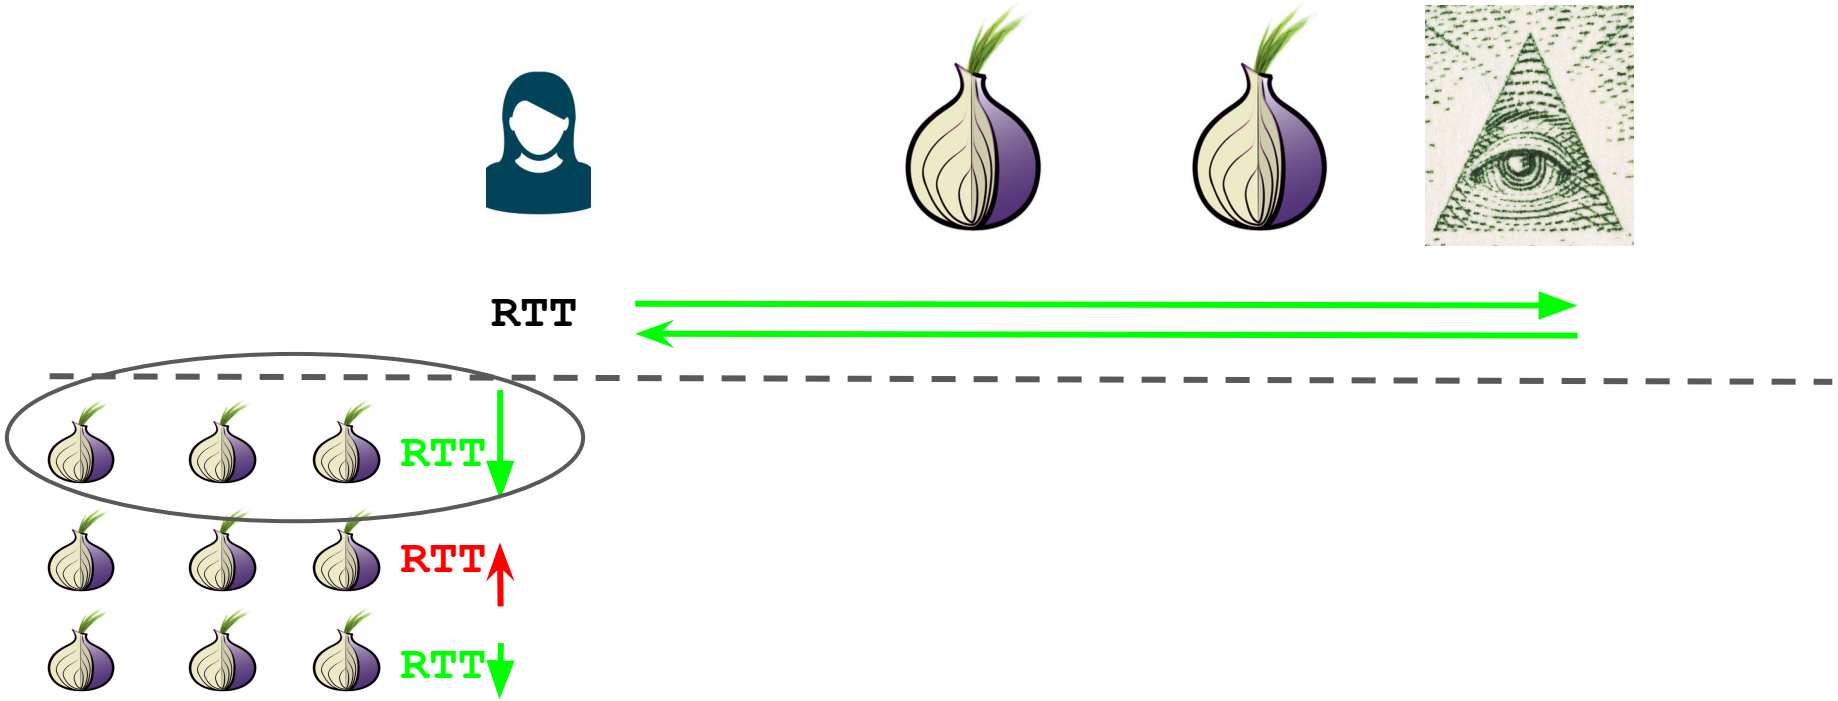


PredicTor

PredicTor

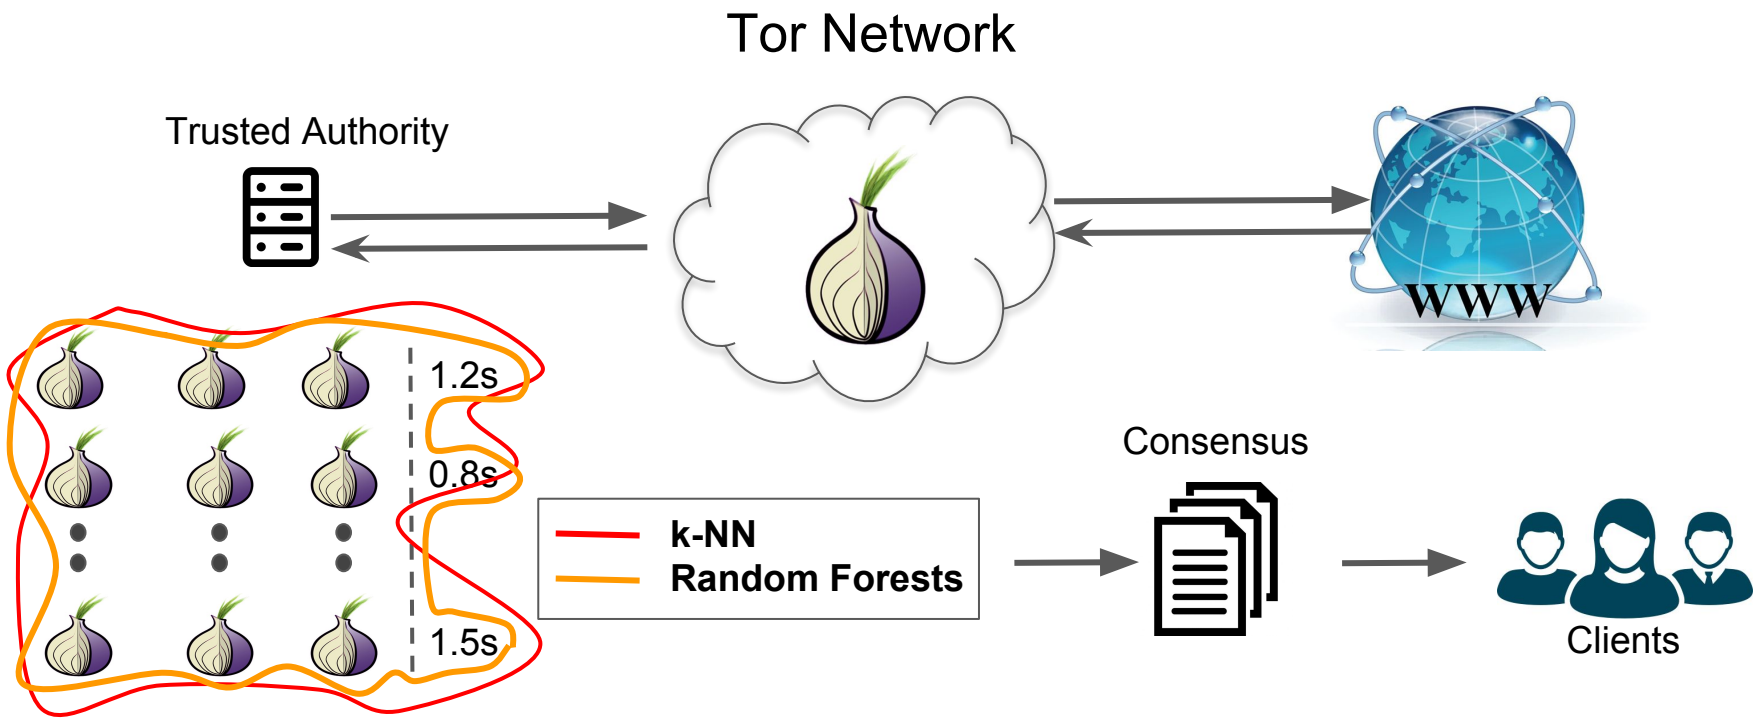
Tor Network
Trusted Authority
WWW
1.2s
0.8s
1.5s
k-NN
Random Forests
Consensus
Clients

Feature Extraction

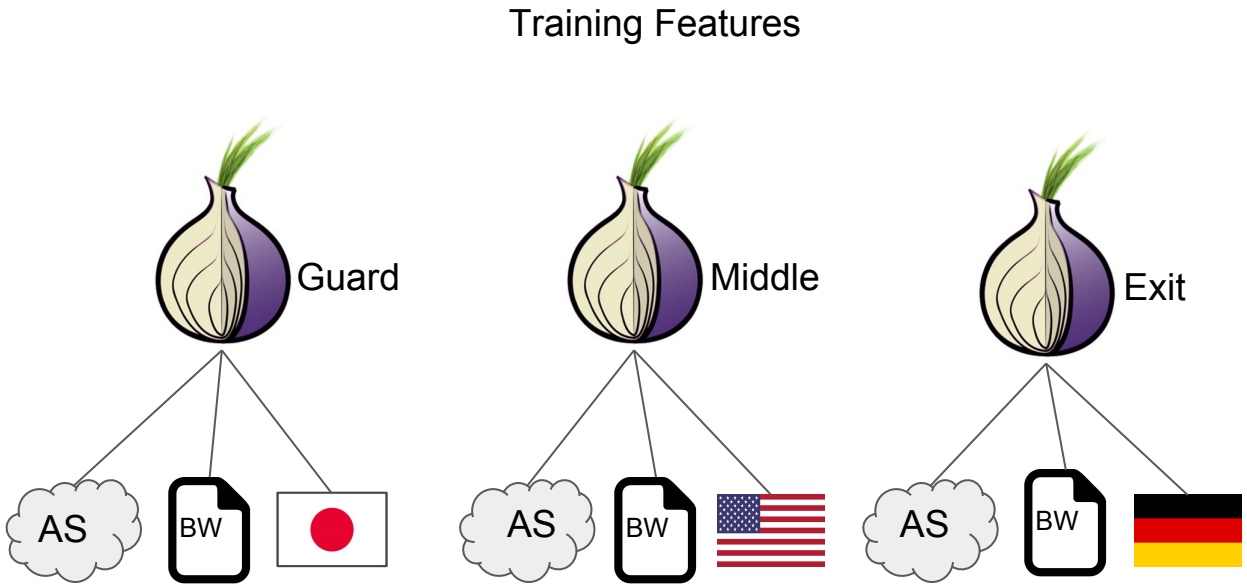

Evaluating Accuracy

Shadow Simulation

- 1000 clients, 400 relays, 70 servers
- 320 KiB
- Training set: 120,000 streams
- Test set: 25,000 streams

Live Tor

- Server hosting 20 instances of Tor
- 80 KiB from a US server
- Training set: 50,000 streams
- Test set: 20,000 streams

Evaluating Accuracy

Model	Shadow	Live Tor
k-NN	70%	64%
Random Forests	**76%**	**70%**

PredicTor Evaluation

Implemented PredicTor in the Tor source code

- Tested on Shadow and Live Tor
- Compared with
  - BW (Vanilla)
  - Congestion Aware Routing (CAR)
  - Snader and Borisov (SB) - 9
  - SB-15

Shadow Experiment

PredicTor Improved Performance

- 23% compared to Vanilla
- 13% compared to CAR
- Speed up over 500ms in the med.
- Over 1.5s in the 90th.
- SB-9 and SB-15 performed the slowest.

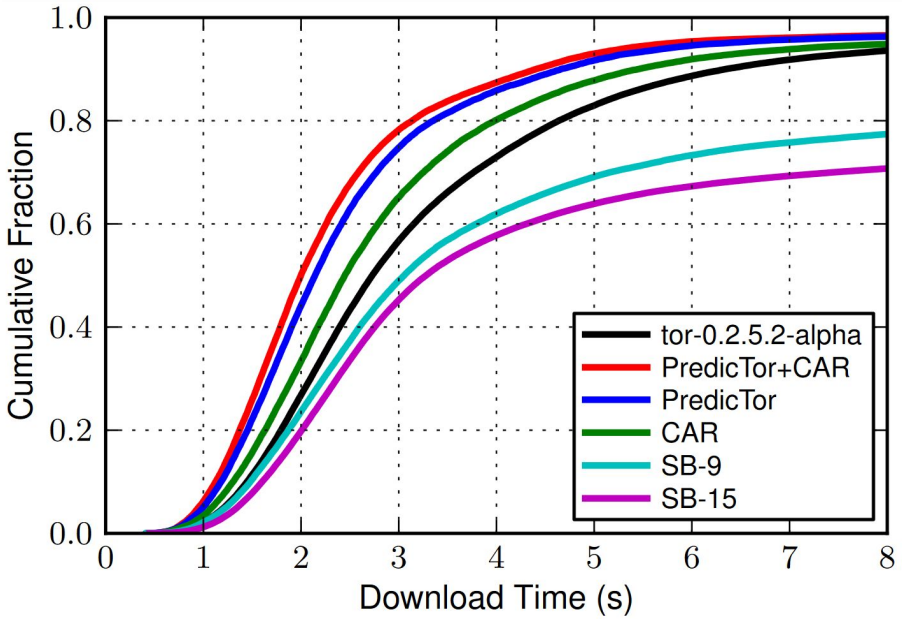

Live Tor Experiment

PredicTor Improved Performance

- 11% compared to Vanilla
- 6% compared to CAR
- Over 1.0s in the 90th.
- SB-9 and SB-15 performed the slowest.

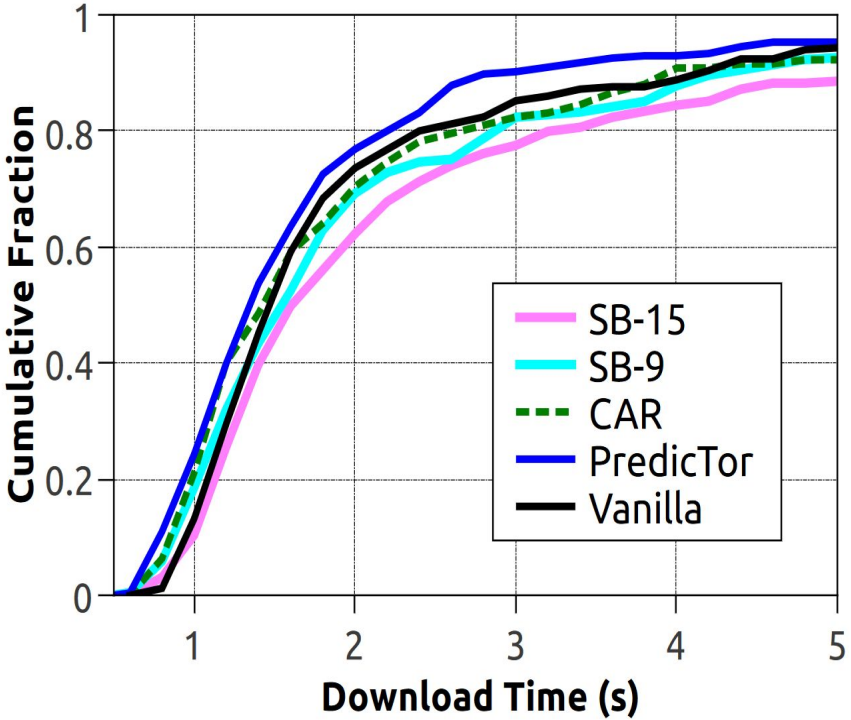

Live Tor Experiment

Circuit Bandwidth

- SB-9
  - 22% BW compared to Vanilla
- SB-15
  - 97% BW compared to Vanilla
- Indicates
  - relays experience persistent congestion
  - performance gains in PredicTor are not solely attributed to BW.

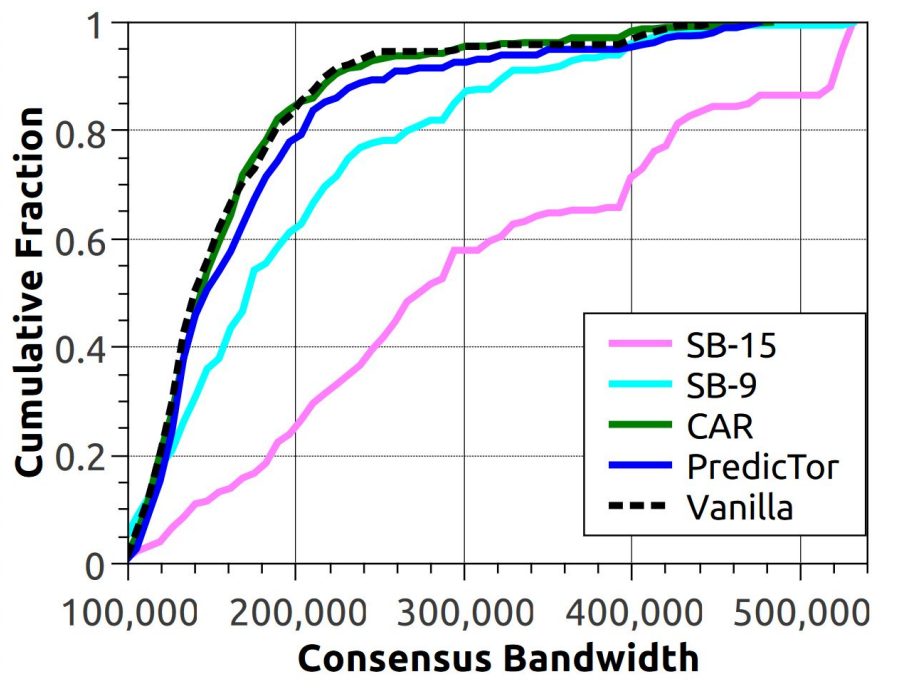

Key Findings

Shadow Simulation

- Relay Utilization
  - SB-9, SB-15 utilized 50%
  - Vanilla utilized 65%
  - PredicTor utilized 85%

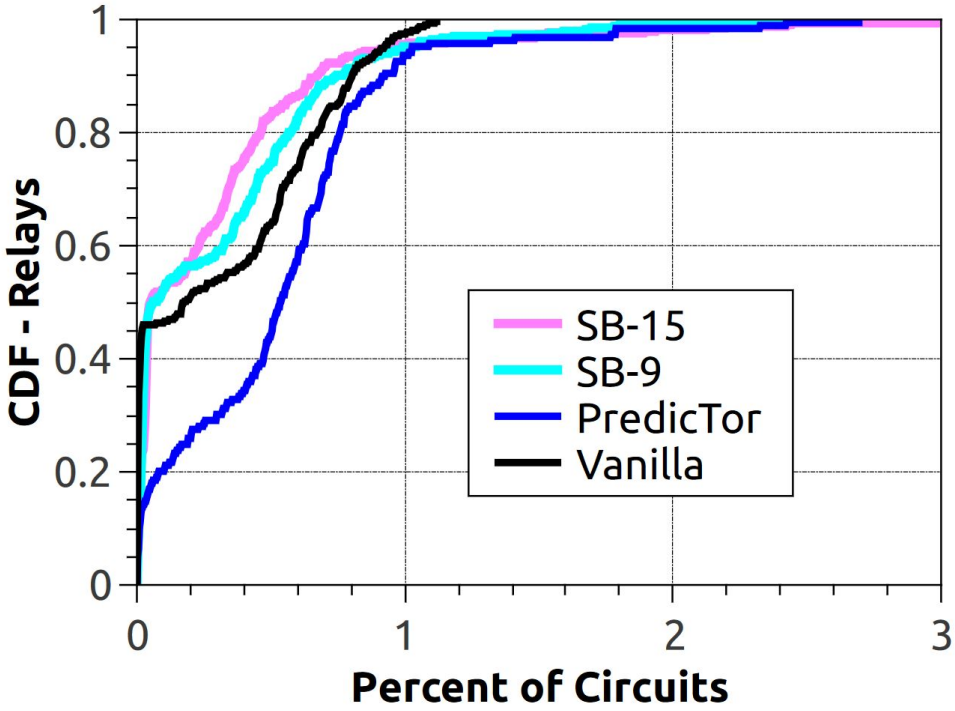

Key Findings

Live Tor Experiment

- Circuit Length
  - 680 km shorter compared to vanilla.

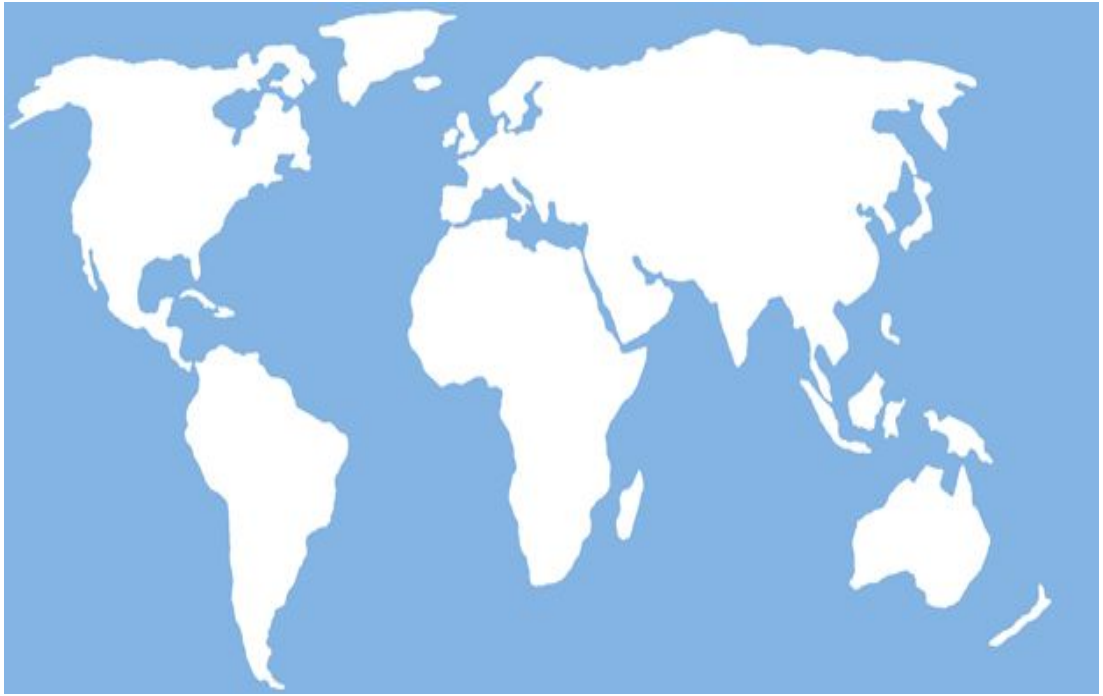

PredicTor Performance Gains

- Avoiding nodes with persistent congestion.
- Better relay utilization.
- Builds circuits of shorter geographic distance.

Security Evaluation

- Entropy based metrics
- All-or-nothing compromise
- AnoA Framework

Client AS Inference (CLASI)

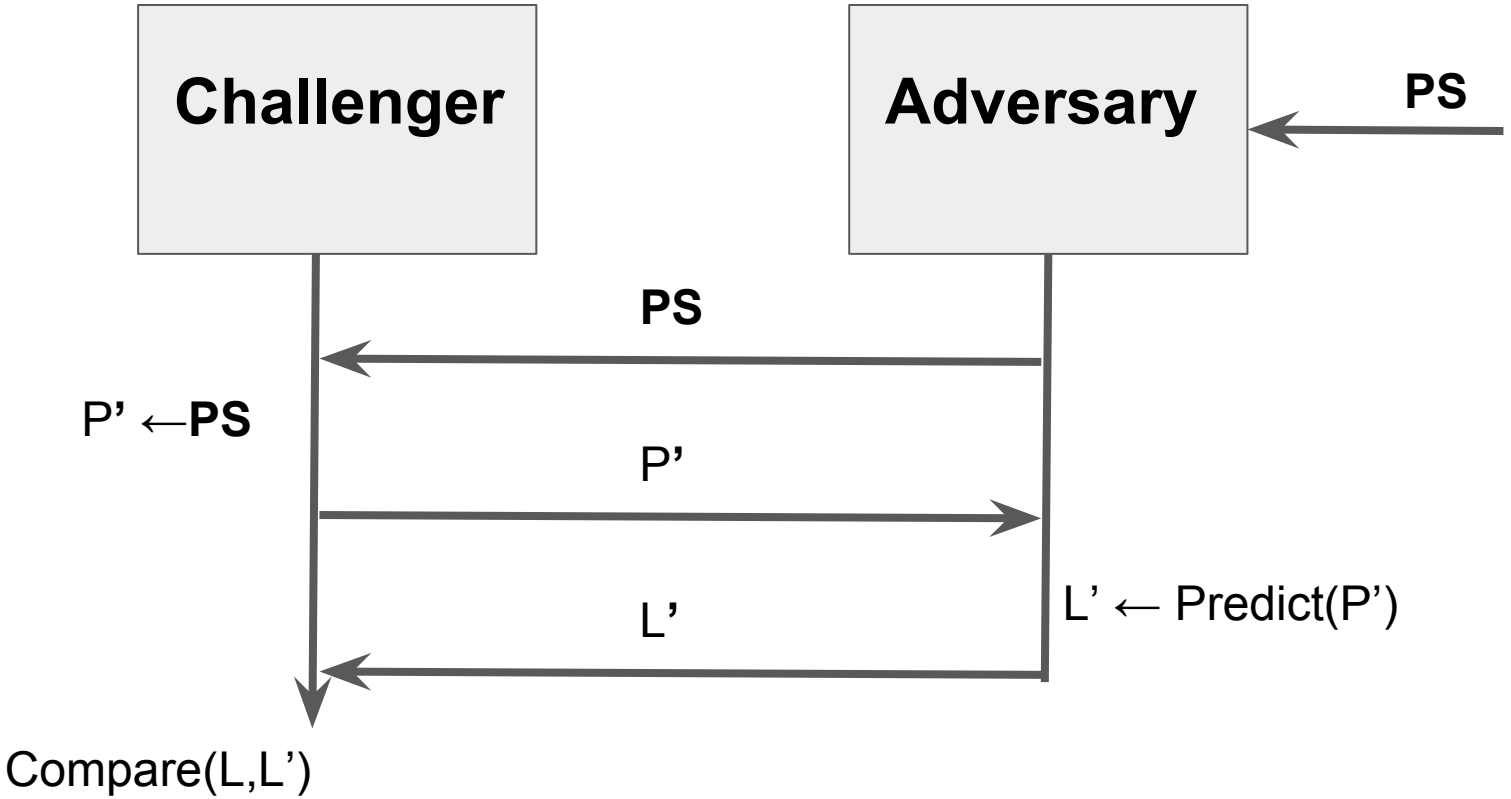

Feature Extraction

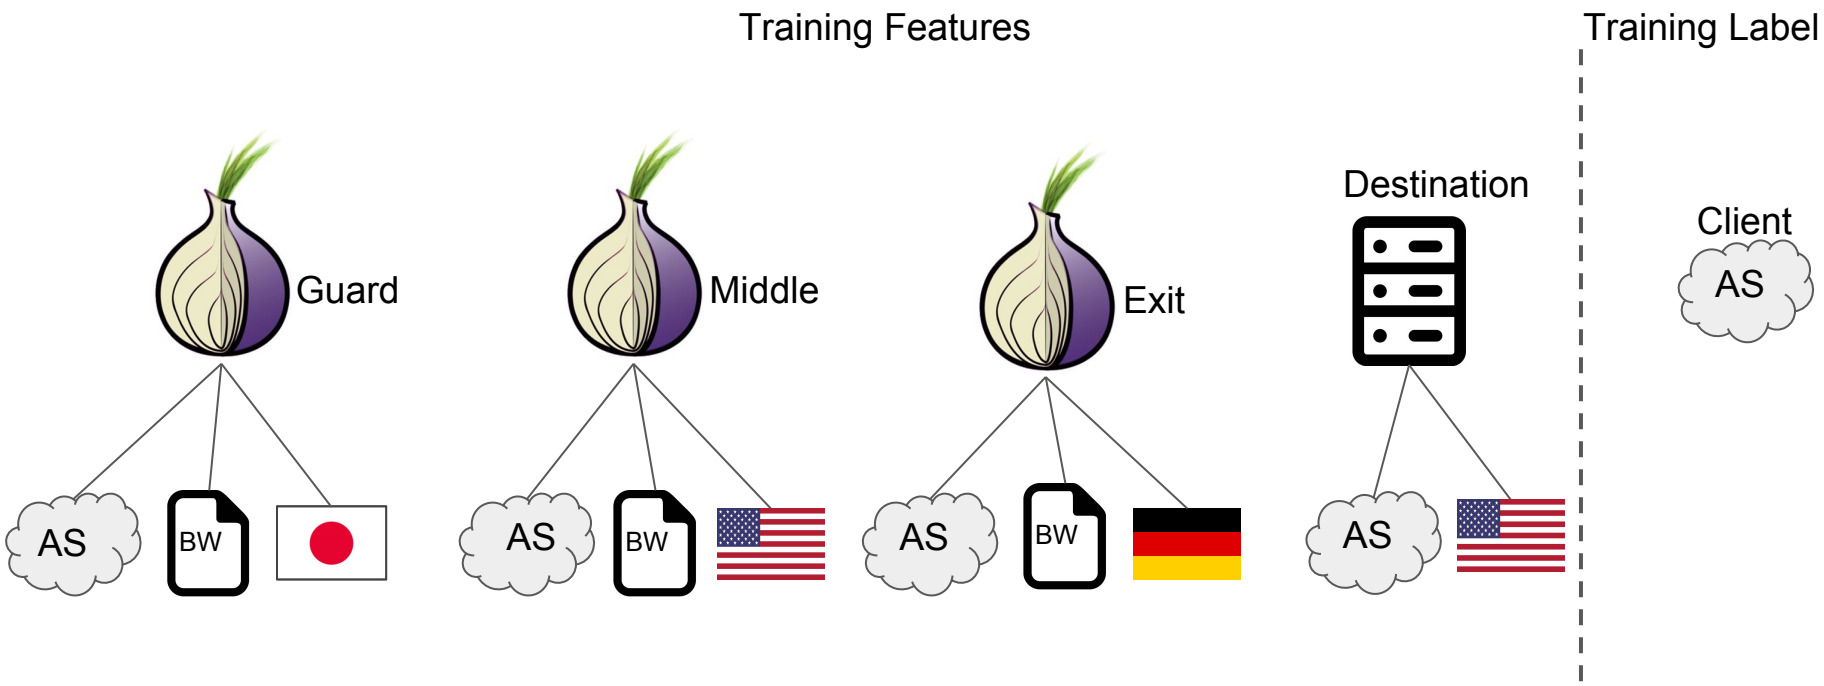

Client AS Inference (CLASI)

$$Pr[L = L'] = \frac{1}{S_L} + \varepsilon_s$$

PredicTor Security Evaluation CLASI

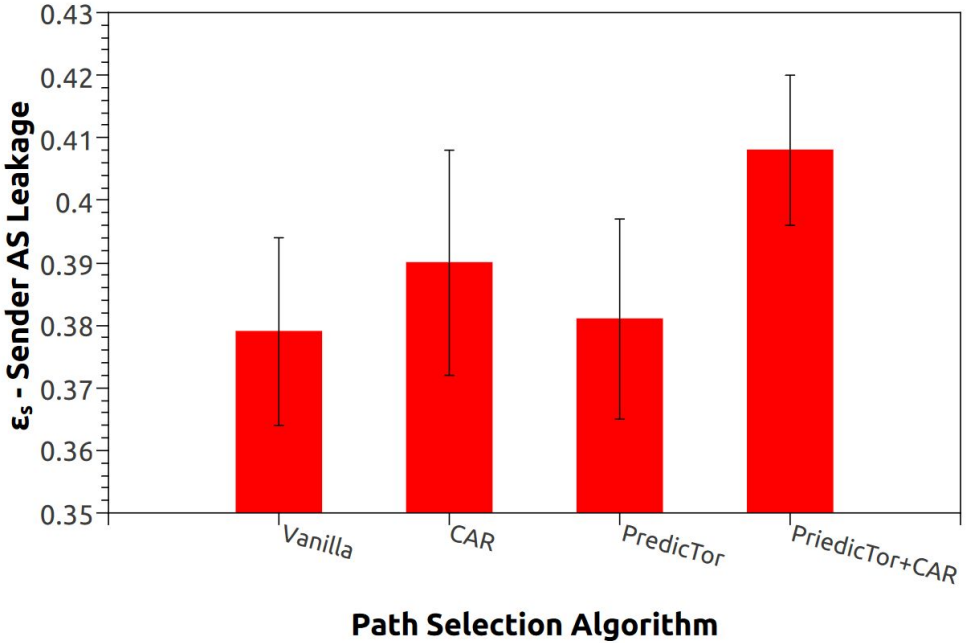

PredicTor Security Evaluation

Uniformity Degree

Algorithm	Uniformity Degree
Vanilla	.84
CAR	.83
PredicTor	.79

Conclusion

PredicTor performance gains

- Avoiding congestion
- Load distribution
- Shorter circuits

PredicTor security evaluation

- PredicTor had Similar sender AS leakage compared to Vanilla
- Lower AS leakage compared to CAR

Conclude: PredicTor had the best security / performance trade-off compared to both Vanilla and CAR.

Questions?